

AD-A193 585

IMPORTANCE SAMPLING FOR STOCHASTIC SIMULATIONS(U)

1/1

STANFORD UNIV CA DEPT OF OPERATIONS RESEARCH

P H GIVIN ET AL AUG 87 TR-24 ARO-20927-28-M

UNCLASSIFIED

DARCT-84-R-0030

F/G 12/3


ML

END

DATE

FILED

78



DTIC FILE COPY

ARO 20927.20-MA

(2)

AD-A193 585

IMPORTANCE SAMPLING FOR STOCHASTIC SIMULATIONS

by

Peter W. Glynn

and

Donald L. Iglehart

TECHNICAL REPORT NO. 24

August 1987

Prepared under the Auspices

of

U.S. Army Research Contract*
DAAG29-84-K-0030

Approved for public release: distribution unlimited.

Reproduction in whole or in part is permitted for any
purpose of the United States government.

DEPARTMENT OF OPERATIONS RESEARCH
STANFORD UNIVERSITY
STANFORD, CALIFORNIA

DTIC
ELECTE

APR 12 1988

*This research was also partially supported under
National Science Foundation Grants DCR 85-09668 and
ECS-8404809.

IMPORTANCE SAMPLING FOR STOCHASTIC SIMULATIONS

by

Peter W. Glynn¹
University of Wisconsin-Madison

and

Donald L. Iglehart²
Stanford University

Accession For	
NTIS CRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution	
Availability Codes	
Dist	ADDITIONAL OF
A-1	Special
DTIC COPY INSPECTED 7	

ABSTRACT

Importance sampling is one of the classical variance reduction techniques for increasing the efficiency of Monte Carlo algorithms for estimating integrals. The basic idea is to replace the original random mechanism in the simulation by a new one and at the same time modify the function being integrated. In this paper the idea is extended to problems arising in the simulation of stochastic systems. Discrete-time Markov chains, continuous-time Markov chains, and generalized semi-Markov processes are covered. Applications are given to a $GI/G/1$ queueing problem and response surface estimation. Computation of the theoretical moments arising in importance sampling is discussed and some numerical examples given.

¹ Supported by NSF Grant ECS-9404809 and by U. S. Army Research Contract No. DAAG29-84-K-0030

² Supported by NSF Grant DCR-85-09668 and by U.S. Army Research Contract No. DAAG29-84-K-0030.

1. INTRODUCTION

The use of importance sampling has long been recognized as a useful technique for increasing the efficiency of Monte Carlo algorithms for numerically evaluating integrals. Given the goal of calculating $I = \int_{-\infty}^{\infty} f(x)dx$, importance sampling involves choosing a probability density $g(\cdot)$ and observing that I can be represented as

$$I = E\{f(X)/g(X)\} \quad (1)$$

where X is a r.v. with density g . One then samples repeatedly from the density g , and estimates I by the sample mean of the observations $f(X)/g(X)$. The term importance sampling derives from the fact that one can choose g to be large in the regions that are most important, namely where $|f|$ is largest. Substantial efficiency increases over conventional Monte Carlo are possible when this technique is properly implemented.

Our purpose in this paper is to illustrate how the method generalizes to stochastic simulations, to develop some of the basic theory relevant to this applications setting, and to briefly describe some of the promising research currently underway in the area. The rest of the paper is organized as follows. As should be clear from the above discussion, the critical element of importance sampling is the choice of the density g . We therefore discuss methods for constructing suitable "densities" in the context of discrete-time Markov chains (Section 2), continuous-time Markov chains (Section 3), and general discrete-event stochastic systems (Section 4). Section 5 gives an abstract overview of the construction, and relates the importance sampling density g to the likelihood ratio of statistics and the Radon-Nikodym derivative of measure-theoretic probability. Section 6 applies the ideas of Section 5 to the processes discussed in Sections 2-4. In Section 7 we discuss the problem of selecting the optimal density g over a suitably defined parametric family of importance sampling densities. An application of this idea to the $GI/G/1$ queue is given. Section 8 deals with the problem of response surface estimation. Here it is shown that sometimes the simulator can estimate an expectation depending on a parameter by only simulating at one value of the parameter. Finally, in Section 9 we discuss the computation of a number of the moments arising in importance sampling and give a numerical example of the technique.

2. IMPORTANCE SAMPLING FOR DISCRETE-TIME MARKOV CHAINS

Suppose that we are assigned the following discrete-time Markov chain simulation problem:

problem:

Given a transition matrix P , initial distribution μ , and real-valued function f , calculate

$$\alpha = E_P f(X_0, X_1, \dots)$$

where $X = \{X_n : n \geq 0\}$ is the Markov chain associated with P and μ . ($E_P(\cdot)$ denotes the expectation on the path space of X associated with P and μ .)

We will now show how to construct an analogue to the importance sampling density g in this problem context. Suppose we represent μ and P respectively as $\mu = (\mu(i) : i \in S)$ and $P = (P(i, j) : i, j \in S)$, where $S = \{0, 1, \dots\}$ is the state space of X . Then, if $f(X_0, X_1, \dots) = f(X_0, X_1, \dots, X_n)$ (i.e. f depends on the trajectory of X only up to the deterministic time n), it follows that

$$E_P f(X_0, \dots, X_n) = \sum_{i_0, \dots, i_n} f(i_0, \dots, i_n) \mu(i_0) \prod_{k=0}^{n-1} P(i_k, i_{k+1}). \quad (2)$$

Let $g : S^{n+1} \rightarrow \mathbb{R}$ be a probability mass function (p.m.f.) on S^{n+1} with the property that

$$\mu(i_0) \prod_{k=0}^{n-1} P(i_k, i_{k+1}) > 0 \text{ implies } g(i_0, \dots, i_n) > 0. \quad (3)$$

(e.g. $g(i_0, \dots, i_n) = h(i_0) \dots h(i_n)$ with h a strictly positive p.m.f. on S). By multiplying and dividing appropriately in (2) (we need (3) to do this), we obtain

$$\begin{aligned} \alpha &= \sum_{i_0, \dots, i_n} f(i_0, \dots, i_n) \frac{\mu(i_0) \prod_{k=0}^{n-1} P(i_k, i_{k+1})}{g(i_0, \dots, i_n)} g(i_0, \dots, i_n) \\ &= E_g f(X_0, \dots, X_n) \frac{\mu(X_0) \prod_{k=0}^{n-1} P(X_k, X_{k+1})}{g(X_0, \dots, X_n)} \end{aligned} \quad (4)$$

when $E_g(\cdot)$ is the expectation over trajectories X having p.m.f. g ; formula (4) is the discrete-time Markov chain analogue to (1).

By relation (4), α is representable as the expectation of a r.v. U taken with respect to $E_g(\cdot)$. This suggests that α may be estimated by repeatedly sampling the r.v. $\vec{X} = (X_0, \dots, X_n)$ from the p.m.f. g and forming a sample mean of the observations $U_i = U(\vec{X}_i)$ so generated. To analyze the efficiency of the new estimator, consider a budget t representing

the amount of computer time assigned to numerical evaluation of α . Then, if β_i is a r.v. denoting the amount of computer time required to generate U_i ,

$$N(t) = \begin{cases} \max\{n \geq 1 : \beta_1 + \dots + \beta_n \leq t\}, & \beta_1 \geq t \\ 0; & \beta_1 < t \end{cases} \quad (5)$$

is the number of U_i 's generated by time t . Thus, the estimator

$$O(t) = \begin{cases} \frac{1}{N(t)}(U_1 + \dots + U_{N(t)}); & N(t) \geq 1 \\ 0; & N(t) = 0 \end{cases} \quad (6)$$

is the estimator available after t units of computational effort have been expended. By appealing to Section 5 of Glynn and Whitt (1986), the following result is easily proved.

Theorem 1. If $E_g(U_1^2 + \beta_1) < \infty$, then $t^{1/2}(O(t) - \alpha) \Rightarrow \sigma(g)N(0, 1)$ as $t \rightarrow \infty$, where $\sigma^2(g) = E_g\beta_1 \cdot \text{var}_g U_1$.

Note that if we choose

$$\hat{g}(i_0, \dots, i_n) = \mu(i_0) \prod_{k=0}^{n-1} P(i_k, i_{k+1})$$

then the above importance sampling algorithm with $g = \hat{g}$ reduces to the conventional Monte Carlo procedure for calculating α . Thus, we obtain an improvement in efficiency by using importance sampling p.m.f. g provided that $\sigma^2(g) \leq \sigma^2(\hat{g})$.

Much of the previous literature on importance sampling assumes that $E_g\beta_1$ is independent of g ; in this case, efficiency is maximized by reducing the variance $\text{var}_g \beta_1$ to its minimal level. Clearly, the assumption that the expected computational effort $E_g\beta_1$ is insensitive to g is especially unpalatable in this problem setting. In fact, we would typically expect here that $E_g\beta_1$ will be a major component of $\sigma^2(g)$, particularly when g is a joint p.m.f. on S^{n+1} from which it is expensive to generate variates. However, it does seem reasonable to expect that $E_g\beta_1 \approx E_{\hat{g}}\beta_1$ when g is itself a Markov probability in S^{n+1} , by which we mean that there exists an initial distribution ν and transition matrices K_k such that

$$g(i_0, \dots, i_n) = \nu(i_0) \cdot \prod_{k=0}^{n-1} K_k(i_k, i_{k+1}). \quad (7)$$

Of course, in order that g satisfy (3), it is necessary that $\mu(i) > 0$ imply $\nu(i) > 0$ and $P(i, j) > 0$ imply $K_k(i, j) > 0$.

The probability g is said to be time-homogeneous Markov if $K_j = K$ for $j \geq 0$. In this case, $\alpha = E_K f(X_0, \dots, X_n) L_n(P, K)$, where $E_K \equiv E_g$ with g given by (7) and

$$L_n(P, K) = \frac{\mu(X_0)}{\nu(X_0)} \cdot \prod_{k=0}^{n-1} \frac{P(X_k, X_{k+1})}{K(X_k, X_{k+1})}. \quad (8)$$

Note that to generate a r.v. \tilde{X} from the Markov probability g given by (7), we generate X_0 according to ν and recursively generate the remaining X_i 's from the transition matrices K_k . Thus, the effort required to generate \tilde{X} under such a Markov probability should be roughly comparable to that required to generate \tilde{X} under P . However, because of the particular simplicity of generating from time-homogeneous Markov probabilities, we shall restrict our attention to this class of sampling "densities" g for the remainder of this section.

We now turn our attention to the class of estimation problems $\alpha = E_P f(X_0, X_1, \dots)$ for which there exists no integer n such that $f(X_0, X_1, \dots) = f(X_0, X_1, \dots, X_n)$. This occurs, for instance, when we wish to calculate $\alpha = E_P T$, when $T = \inf\{n \geq 0 : X_n \in A\}$; such problems arise, for example, when it is necessary to calculate the expected time until a buffer overflows in a manufacturing system. (Here, A is the set of states corresponding to a buffer overflow.)

Suppose that $f(X_0, X_1, \dots) = f(X_0, \dots, X_T)$ where T is a stopping time for the process X . (A stopping time is a random time with the property that the event $\{T = n\}$ is determined completely by the history of (X_0, \dots, X_n) ; see p. 118 of Çinlar (1975) for details.) Our next proposition shows that (8) extends in the natural way to estimating expectations of this form.

Proposition 1. Suppose that $E_P |f(X_0, \dots, X_T)| < \infty$ with T a stopping time. Then

$$E_P f(X_0, \dots, X_T) = E_K f(X_0, \dots, X_T) L_T(P, K)$$

where

$$L_T(P, K) = \frac{\mu(X_0)}{\nu(X_0)} \prod_{k=0}^{T-1} \frac{P(X_k, X_{k+1})}{K(X_k, X_{k+1})}.$$

Proposition 1 shows that importance sampling extends easily to the case where f depends on X up to a stopping time T . Efficiency increases are obtained by appropriately choosing ν and K , and then simulating X under initial distribution ν and transition matrix K ; Proposition 1 is then used to construct the sample mean estimator.

We conclude this section with a description of how importance sampling extends to steady-state estimation. We recall that the steady-state estimation problem requires calculating $\alpha = E_P f(X_0, X_1, \dots)$ where

$$f(X_0, X_1, \dots) = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} h(X_k). \quad (9)$$

i.e. $f(X_0, X_1, \dots)$ is the long-run average of $\{h(X_n) : n \geq 0\}$ for some function h . A naive approach to calculating α would involve extending Proposition 1 to functions f defined through (9). In other words, one might hope that

$$\alpha = E_K \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} h(X_k) \cdot L_\infty(P, K), \quad (10)$$

where $L_\infty(P, K) = \lim_{n \rightarrow \infty} L_n(P, K)$. However, it turns out that (10) is false in general, and so we conclude that importance sampling can break down for functionals f which depend on the infinite history of the process $X = \{X_n : n \geq 0\}$. Since this observation is an important one, we offer an elementary demonstration that (10) is, in general, false. (In the language of measure-theoretic probability, (10) breaks down because the measures on the infinite path space of X associated with P and K are mutually singular; see Section 5 for more details.)

Suppose that h is strictly positive and that X is an irreducible Markov chain on finite state space S under transition matrix P . Clearly, α is then positive. We shall see, however, that the right-hand side of (10) vanishes, unless $P = K$. To prove this, it suffices to show that $L_\infty(P, K) = 0$ under $P_K(\cdot)$. ($P_K(\cdot)$ is the probability corresponding to $E_K(\cdot)$.) Clearly, the result is immediate if $P(i, j)$ vanishes when $K(i, j)$ is positive, since the irreducibility and finite state space guarantees that (X_n, X_{n+1}) will eventually visit such an (i, j) under $P_K(\cdot)$. Otherwise, observe that

$$L_\infty(P, K) = \frac{\mu(X_0)}{\gamma(X_0)} \lim_{n \rightarrow \infty} \exp \left[\sum_{k=0}^{n-1} \phi(X_k, X_{k+1}) \right] \quad (11)$$

when $\phi(i, j) = \log(P(i, j)/K(i, j))$. But

$$\frac{1}{n} \sum_{k=0}^{n-1} \phi(X_k, X_{k+1}) \rightarrow \sum_{i, j \in S} p(i) K(i, j) \phi(i, j) \quad (12)$$

P_K a.s., where $p = (p(i) : i \in S)$ are the stationary probabilities of K . By the strict concavity of $\log(\cdot)$,

$$\sum_{i, j \in S} p(i) K(i, j) \log \left(\frac{P(i, j)}{K(i, j)} \right) < \log \left(\sum_{i, j} p(i) P(i, j) \right) = 0$$

with strict inequality when the points $P(i, j)/K(i, j)$ are distinct (i.e. $P(i, j) \neq K(i, j)$). By (12), $\sum_{k=0}^{n-1} \phi(X_k, X_{k+1}) \rightarrow -\infty$ and thus $L_\infty(P, K) = 0$ by (11).

In spite of the above difficulty, it turns out that importance sampling can be applied to the steady-state estimation problem. The idea is to use the regenerative structure

of X to reduce the infinite-horizon steady-state behavior to a finite-horizon analysis over regenerative cycles. To be precise, for a given (but arbitrary) state $i \in S$, let $T(i) = \inf\{n \geq 1 : X_n = i\}$, and set

$$\begin{aligned} Y &= \sum_{k=0}^{T(i)-1} h(X_k) \\ \hat{Y} &= \sum_{k=0}^{T(i)-1} |h(X_k)| \\ \tau &= T(i). \end{aligned}$$

It is well known that the steady-state mean of h can be expressed as $\alpha = E_P\{Y|X_0 = i\} / E_P\{\tau|X_0 = i\}$. By observing that T is a stopping time and applying Proposition 1, we obtain the following result.

Proposition 2. Suppose X is a positive recurrent irreducible Markov chain under transition matrix P . If $E_P\{\hat{Y}|X_0 = i\} < \infty$, then

$$E_P \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} h(X_k) = \frac{E_K\{Y L_\tau(P, K)|X_0 = i\}}{E_K\{\tau L_\tau(P, K)|X_0 = i\}}.$$

A second representation for the steady-state mean α can also be derived. Observe that Y is expressed as a sum, and note that

$$E_K \left\{ \sum_{\ell=0}^{n-1} h(X_\ell) L_n(P, K) \right\} = \sum_{\ell=0}^{n-1} E_K \left\{ h(X_\ell) \frac{\mu(X_0)}{\gamma(X_0)} \cdot \prod_{j=0}^{n-1} \frac{P(X_j, X_{j+1})}{K(X_j, X_{j+1})} \right\}.$$

It is easily verified that

$$E_K \left\{ \prod_{j=\ell}^n \frac{P(X_j, X_{j+1})}{K(X_j, X_{j+1})} | X_0, \dots, X_\ell \right\} = 1$$

and hence, it follows by conditioning on X_0, \dots, X_ℓ that

$$E_K \left\{ \sum_{\ell=0}^{n-1} h(X_\ell) L_n(P, K) \right\} = E_K \left\{ \sum_{\ell=0}^{n-1} h(X_\ell) L_\ell(P, K) \right\}.$$

The above argument easily extends, as in the proof of Proposition 1, to stopping times T . Applying the stopping time version to the regenerative representation of the steady-state mean yields the following result.

Proposition 3. Suppose X is a positive recurrent irreducible Markov chain under transition matrix P . If $E_P\{\hat{Y}|X_0 = i\} < \infty$, then

$$E_P \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} h(X_k) = \frac{E_K \left\{ \sum_{\ell=0}^{\tau-1} h(X_\ell) L_\ell(P, K) | X_0 = i \right\}}{E_K \left\{ \sum_{\ell=0}^{\tau-1} L_\ell(P, K) | X_0 = i \right\}}$$

As a consequence of Propositions 2 and 3, we see that the steady-state mean α can be expressed as the ratio of two expectations; standard methods can then be applied to estimate α . In particular, suppose $\{(YL)_n, (\tau L)_n : n \geq 1\}$ is a sequence of i.i.d. replicates of $(YL_r(P, K), \tau L_r(P, K))$ generated under $P_K(\cdot)$. Then, the estimator $r_1(t)$ for α is the ratio estimator, available after t time units of computation have been expended, of the sample mean of $(YL)_n$'s to $(\tau L)_n$'s. Similarly, the ratio estimator $r_2(t)$ can be constructed from Proposition 3. Let $\beta(1), \beta(2)$ be the computation time required to generate the random vectors $(YL_r(P, K), \tau L_r(P, K)), (\sum_{\ell=0}^{r-1} h(X_\ell) L_\ell(P, K), \sum_{\ell=0}^{r-1} L_\ell(P, K))$, respectively, under $P_K(\cdot)$. The following result, which follows immediately from Section 5 of Glynn and Whitt (1986), describes the rates of convergence of the two steady-state estimators.

Theorem 2. Suppose X is a positive recurrent irreducible Markov chain under transition matrix P . If $E_K\{(\hat{Y}^2 + \tau^2)L_r^2|X_0 = i\} < \infty$ and $E_K(\beta(1) + \beta(2)) < \infty$, then

$$t^{1/2}(r_i(t) - \alpha) \Rightarrow \sigma_i N(0, 1)$$

as $t \rightarrow \infty$, for $i = 1, 2$, where

$$\begin{aligned} \sigma_1^2 &= E_K\{\beta(1)|X_0 = i\} \cdot E_K\{(Y - \alpha\tau)^2 L_r^2(P, K)|X_0 = i\} / (E_K\{\tau L_r(P, K)|X_0 = i\})^2 \\ \sigma_2^2 &= E_K\{\beta(2)|X_0 = i\} \cdot \frac{E_K\left\{\left(\sum_{\ell=0}^{r-1} (h(X_\ell) - \alpha)L_\ell(P, K)\right)^2 |X_0 = i\right\}}{\left(E_K\left\{\sum_{\ell=0}^{r-1} L_\ell(P, K)|X_0 = i\right\}\right)^2}. \end{aligned}$$

As in Theorem 1, confidence intervals for α follow easily from this central limit theorem. Theorem 2 raises the question of which of the two steady-state estimators is more efficient. Note that the proofs of Proposition 2 and 3 show that

$$E_K\{\tau L_r(P, K)|X_0 = i\} = E_K\left\{\sum_{\ell=0}^{r-1} L_\ell(P, K)|X_0 = i\right\} = E_P\{\tau|X_0 = i\}.$$

Thus, if $\beta(1)$ is approximately equal to $\beta(2)$ (which seems reasonable in most applications), $r_2(t)$ is more efficient than $r_1(t)$ if (and only if)

$$E_K\left\{\left(\sum_{\ell=0}^{r-1} (h(X_\ell) - \alpha)L_\ell(P, K)\right)^2 |X_0 = i\right\} \leq E_K\left\{\left(\sum_{k=0}^{r-1} (h(X_k) - \alpha)\right)^2 L_r^2(P, K)|X_0 = i\right\}. \quad (13)$$

The choice of return state i is another natural question which arises in the regenerative context. If one assumes that the computational effort equals τ (this is a common assumption in regenerative analysis), then it is well known (see Crane and Iglehart (1975)) that

the choice of return state does not affect the efficiency parameter σ^2 when conventional regenerative output analysis is used. On the other hand, the choice of return state i does have an effect on the asymptotic efficiency in the importance sampling case currently under study here; an example illustrating this appears in Section 9. Although we have no mathematical theory to guide the choice of return state, it seems reasonable to employ the heuristic of choosing that state i with the shortest expected return time under $P_K(\cdot)$.

3. IMPORTANCE SAMPLING FOR CONTINUOUS-TIME MARKOV CHAINS

In this section, we show how to extend the ideas of Section 2 to continuous-time Markov chains. To be precise, given a non-explosive continuous-time Markov chain $X = \{X(t) : t \geq 0\}$ with (conservative) generator Q and initial distribution μ , our goal is to apply importance sampling to the estimation of

$$\alpha = E_Q f(X),$$

where f is real-valued. ($E_Q(\cdot)$ corresponds to the expectation on the path-space of X associated with using generator Q , when initialized with μ).

As in Section 2, it is possible to do importance sampling in which the sampling distribution does not correspond to a continuous-time Markov chain. However, as argued in Section 2, it seems reasonable to restrict ourselves to using sampling distributions which are themselves Markov, and we shall do so here.

The simulation of a continuous-time Markov chain generally proceeds on the time scale of state transitions. By this, we mean that any event-driven algorithm for generating the sample path of such a process will evolve iteratively by determining the "next" state and corresponding holding time for that state. Thus, after n iterations of the algorithm, we have simulated the process up to time S_n^- , where S_n is the instant at which the chain enters the $(n+1)$ 'st state. In other words, the natural time scale for discrete-event simulation of such a chain is the time scale corresponding to $\{S_n : n \geq 0\}$. As a consequence, we shall initially concentrate on studying estimation problems for which $f(X)$ can be expressed as $f(X(t) : 0 \leq t < S_n)$ for some $n \geq 0$. Note that such a function f can be re-expressed as $\tilde{f}(U_0, V_0, U_1, V_1, \dots, U_n, V_n)$, where U_i is the i 'th state visited by the chain and V_i is the time spent in that state.

Recall that $\{U_n : n \geq 0\}$ is a discrete-time Markov chain with transition matrix R defined by $R(i, j) = Q(i, j)/q(i)$ for $i \neq j$ and $R(i, i) = 0$, where $q(i) = -Q(i, i)$. Furthermore, conditional on $\{U_n : n \geq 0\}$, the V_n 's are independent exponential r.v.'s for which the (conditional) mean of V_n is $1/q(U_n)$. From these observations, it is immediate that

$$\begin{aligned} \alpha &= E_Q f(X(t) : 0 \leq t < S_n) \\ &= E_Q \tilde{f}(U_0, V_0, \dots, U_n, V_n) \\ &= \sum_{u_0, \dots, u_n} \int_0^\infty \dots \int_0^\infty \tilde{f}(u_0, v_0, \dots, u_n, v_n) \mu(u_0) \cdot \prod_{j=0}^{n-1} R(u_j, u_{j+1}) \prod_{j=0}^n q(u_j) \exp(-q(u_j)v_j) dv_0 \dots dv_n \\ &= \sum_{u_0, \dots, u_n} \int_0^\infty \dots \int_0^\infty \tilde{f}(u_0, v_0, \dots, u_n, v_n) \mu(u_0) \prod_{j=0}^{n-1} Q(u_j, u_{j+1}) \exp\left(-\sum_{j=0}^n q(u_j)v_j\right) dv_0, \dots, dv_n. \end{aligned} \quad (14)$$

Suppose that we wish to estimate α by simulating a chain with a different initial distribution γ and/or generator A . We assume that A is non-explosive and conservative and that the analogue of condition (3) holds:

$$\begin{aligned} \mu(i) > 0 &\text{ implies } \gamma(i) > 0 \text{ and} \\ Q(i, j) \neq 0 &\text{ implies } A(i, j) \neq 0. \end{aligned} \quad (15)$$

From (14), it is evident that an appropriate multiplication and division yields

$$\begin{aligned} \alpha &= E_A \tilde{f}(U_0, V_0, \dots, U_n, V_n) L_n(Q, A) \\ &= E_A f(X(t) : 0 \leq t < S_n) L_n(Q, A) \end{aligned}$$

where

$$L_n(Q, A) = \frac{\mu(U_0)}{\gamma(U_0)} \prod_{j=0}^{n-1} \frac{Q(U_j, U_{j+1})}{A(U_j, U_{j+1})} \exp\left(\int_0^{S_n} (a(X(s)) - q(X(s))) ds\right) \quad (16)$$

and $a(i) = -A(i, i)$. Hence, α can be estimated via importance sampling. We simulate the chain under generator A , as initialized by γ , up to time S_n . At the end of the run, the quantity $f(X(t) : 0 \leq t < S_n) L_n(Q, A)$ is calculated. By independently replicating these runs many times, we can estimate α by a sample mean of the $f(X(t) : 0 \leq t < S_n) L_n(Q, A)$'s so obtained.

As in Section 2, we now wish to extend our importance sampling methodology to more general functions $f(X)$. Suppose that T is a stopping time relative to the sequence $\{(U_k, V_k) : k \geq 0\}$, by which we mean that the event $\{T = n\}$ is completely determined by $(U_0, V_0), \dots, (U_n, V_n)$. The next result follows from (16) in a manner similar to the way in which Proposition 1 followed from (8).

Theorem 3. If $E_Q|f(X(t) : 0 \leq t < S_T)| < \infty$ with T a (finite-valued) stopping time, then

$$E_Q f(X(t) : 0 \leq t < S_T) = E_A f(X(t) : 0 \leq t < S_T) L_T(Q, A)$$

where

$$L_T(Q, A) = \frac{\mu(U_0)}{\gamma(U_0)} \prod_{j=0}^{T-1} \frac{Q(U_j, U_{j+1})}{A(U_j, U_{j+1})} \exp \left(\int_0^{S_T} (a(X(s)) - q(X(s))) ds \right).$$

An important application of this stopping time result occurs when one is required to estimate the transient behavior of the chain X up to a deterministic time w . In other words, suppose that $f(X) = f(X(t) : 0 \leq t \leq w)$, where w is deterministic. To simulate the process X up to time w requires $T(w)$ iterations of the event-scheduling algorithm, where $T(w) = \inf\{n \geq 0 : S_0 + \dots + S_n \geq w\}$. Clearly, $f(X(t) : 0 \leq t \leq w) = \tilde{f}(X(t) : 0 < t \leq S_{T(w)})$ w.p.l. (Equality may not hold if $S_{T(w)} = w$, but this happens with zero probability.), and hence

$$E_Q f(X(t) : 0 \leq t \leq w) = E_A f(X(t) : 0 \leq t \leq w) L_{T(w)}(Q, A).$$

Formulas similar to this appear on p. 166 of Brémaud (1981), where they were derived using martingale representations of continuous-time Markov chains.

We turn now to the steady-state analysis of continuous-time chains. As in the discrete-time setting of Section 2, we exploit regenerative structure. For a given (arbitrary) state in S , let $\tau = \inf\{n \geq 1 : U_n = i\}$,

$$Y = \int_0^{S_\tau} h(X(s)) ds$$

$$\hat{Y} = \int_0^{S_\tau} |h(X(s))| ds.$$

An analysis similar to that used to obtain Proposition 2 shows that if X is irreducible positive recurrent and if $E_Q\{\hat{Y} | X(0) = i\} < \infty$, then

$$E_Q \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t h(X(s)) ds = \frac{E_A\{Y L_\tau(Q, A) | X(0) = i\}}{E_A\{S_\tau L_\tau(Q, A) | X(0) = i\}}. \quad (17)$$

Although (17) could be used directly to obtain the steady-state means, such an algorithm would be inefficient. In particular, one can apply conditional Monte Carlo, in the form of discrete-time conversion (see also Fox and Glynn (1986)), to (17), thereby yielding the expression

$$E_Q \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t h(X(s)) ds = \frac{E_A \left\{ \sum_{k=0}^{\tau-1} \frac{h(U_k)}{q(U_k)} \prod_{j=0}^{\tau-1} \frac{R(U_j, U_{j+1})}{R'(U_j, U_{j+1})} | X(0) = i \right\}}{E_A \left\{ \sum_{k=0}^{\tau-1} \frac{1}{q(U_k)} \prod_{j=0}^{\tau-1} \frac{R(U_j, U_{j+1})}{R'(U_j, U_{j+1})} | X(0) = i \right\}} \quad (18)$$

where $R'(i, j) = A(i, j)/a(i)$ for $i \neq j$ and $R'(i, i) = 0$. The principle of conditional Monte Carlo guarantees that a ratio estimator based on (18) will have a lower variance than one based on (17); a central limit theorem for this estimator similar to that appearing in Theorem 2 can easily be derived. An alternative derivation of (18) rests on the well-known fact that

$$E_Q \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t h(X(s)) ds = \frac{E_R \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} \frac{h(U_k)}{q(U_k)}}{E_R \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} \frac{1}{q(U_k)}}$$

Proposition 2 can then be applied to the numerator and denominator of the right-hand side to obtain (18).

To conclude this section, we note that the r.v.'s appearing on the right-hand side depend only on the embedded discrete-time chain $\{U_n : n \geq 0\}$; this chain can be simulated under $P_A(\cdot)$ by using the transition matrix R' . Thus, if (18) is used to estimate the steady-state mean, no exponential variates need be generated.

4. IMPORTANCE SAMPLING FOR GENERALIZED SEMI-MARKOV PROCESSES

In the preceding two sections, we have shown how importance sampling applies to discrete and continuous-time Markov chains. We shall now show how importance sampling applies to general discrete-event simulations. As in Section 3, we shall restrict attention to sampling distributions which themselves correspond to discrete-event simulations.

A mathematical framework is convenient for the purposes of this discussion. In particular, we shall model a discrete-event system as a generalized semi-Markov process (GSMP). As described in Glynn (1983), a GSMP is basically a mathematical formalization of a discrete-event simulation; the process is driven by an event-scheduling algorithm, in exactly the same way as is a discrete-event system.

To be precise, let S be a subset of the non-negative integers representing the state space of the GSMP. The GSMP also requires specification of a second finite integer-valued set E ; E corresponds to the set of all possible events that can initiate a state transition. For example, in the single-server queue, E consists of two events, one each for the arrival and departure processes.

To understand the probabilistic dynamics of a GSMP, recall that in order to simulate a discrete-event system, we need to schedule each event possible in the state currently

occupied; we can view the event schedule as being represented by a set of clocks. So, for each event $e \in E(s)$ (the set of events active in state $s \in S$), suppose that $r(s, e)$ is the non-negative rate at which the clock corresponding to event e runs down in state s ; we assume that $\max\{r(s, e) : e \in E(s)\} > 0$ for each state $s \in S$ (i.e. at least one clock has a positive speed) and $\max\{r(s, e) : e \in E(s), s \in S\} < \infty$. Under this assumption, a clock will eventually run down to zero, thereby initiating a state transition. If two or more clocks run down to zero simultaneously, the clock with the minimal index e is regarded as the "trigger" event. The next state s' is then chosen with probability $p(s'; s, e^*)$, where s was the previous state and event $e^* \in E(s)$ was the event that triggered the transition there. In state s' , the clocks in state s that are still active in s' continue to run down (now according to rates $r(s', e)$); the set of these "old" clocks is the set $O(s', s, e^*) = E(s') \cap (E(s) - \{e^*\})$. In general, when state s' is entered, new events (represented by the set $N(s', s, e^*) = E(s') - O(s', s, e^*)$) need to be scheduled. The clock on event $e \in N(s', s, e^*)$ is set (independently of the past) according to a probability distribution $F(\cdot; s', e, s, e^*)$; we require that $F(0; s', e, s, e^*) = 0$ (clocks are set at strictly positive values). The clocks in $E(s')$ are then allowed to run down, and another state transition occurs; the above process is then repeated indefinitely. The state occupied at time t will be denoted as $X(t)$; the process $\{X(t) : t \geq 0\}$ is then the desired GSMP. It should be clear that the above description of a GSMP closely mimics the evolution of a discrete-event system. For more details on the specification of a GSMP, see Burman (1981) or Whitt (1980).

As in Section 3, it is convenient to view the GSMP on a time-scale which is natural to the simulator, namely $\{S_n : n \geq 0\}$, where S_n is the instant at which the GSMP enters the $(n+1)$ 'st state visited. Let $\{(A_n, C_n) : n \geq 0\}$ be the sequence in which A_n represents the n 'th state visited by the GSMP and $C_n = (C_n(e) : e \in E)$ is the vector of clock readings at the instant of entry into state A_n . (If a clock is inactive in state A_n , the clock is set to $+\infty$.) Clearly, the trigger event e_n^* in A_n is a function of A_n and C_n namely

$$e_n^* = \min\{e' \in E : C_n(e')/r(A_n, e') = \min\{C_n(e)/r(A_n, e) : e \in E\}\}.$$

Note that $S_n = \sum_{k=0}^n C_k(e_k^*)/r(A_k, e_k^*)$. A little thought then shows that if $f(X(t) : t \geq 0) = f(X(t) : 0 \leq t < S_n)$, f can be represented as $\tilde{f}(A_0, C_0, \dots, A_n, C_n)$. A further simplification is possible, however. Observe that $C_n = c_n(A_0, C_0, \dots, A_n, C_n)$ (with $c_n(\cdot)$ deterministic), where $C_i = (C_i(e) : e \in E)$ is the vector of "new" clock readings generated at the n 'th transition,

namely

$$\underline{C}_0 = \begin{cases} C_0(e), & e \in E(A_0) \\ +\infty, & e \notin E(A_0) \end{cases}$$

and

$$\underline{C}_n = \begin{cases} C_n(e), & e \in N(A_n, A_{n-1}, e_{n-1}^*) \\ +\infty, & e \notin N(A_n, A_{n-1}, e_{n-1}^*) \end{cases}$$

for $n \geq 1$. We conclude that there exists a (deterministic) function f such that $f(X(t) : 0 \leq t < S_n) = f(A_0, \underline{C}_0, \dots, A_n, \underline{C}_n)$. This representation of f will now be used to develop importance sampling methodology for a GSMP estimation problem of the form

$$\alpha = E_F f(X(t) : 0 \leq t < S_n). \quad (19)$$

($E_F(\cdot)$ is the expectation on the path space of the GSMP associated with transition probability $p(s'; s, e^*)$, clock-setting distributions $F(\cdot; s', e, s, e^*)$, and initial distribution μ ; we assume that the initial distribution takes the form

$$P\{A_0 = s, C_0(e) \in dx_e, e \in E(s)\} = \mu(s) \cdot \prod_{e \in E(s)} F(dx_e, e).$$

As one might expect from our discussion of discrete and continuous-time chains, an analogue to conditions (3) and (15) will be needed. Suppose that we wish to use a sampling distribution which corresponds to a GSMP having identical state space S , event space E , rates $r(s, e)$, but possessing (possibly) different transition probabilities $q(s', e, e^*)$, clock-setting distributions $H(\cdot; s', e, s, e^*)$ and initial distribution γ (γ is characterized by initial state probability $\gamma(s)$ and initial clock-setting distributions $H(dx_e, e)$.) The data of the new GSMP must satisfy:

- 1.) $p(s'; s, e^*) > 0$ implies $q(s'; s, e^*) > 0$ and $\mu(s) > 0$ implies $\gamma(s) > 0$.
- 2.) there exists functions $k(\cdot; s', e, s, e^*)$ and $k(\cdot, e)$ such that

$$F(dx; s', e, s, e^*) = k(x; s', e, s, e^*) H(dx; s', e, s, e^*)$$

$$F(dx, e) = k(x, e) H(dx, e).$$

Since the first condition is self-explanatory, we concentrate on explaining the second condition above. It basically says that wherever H assigns no mass (i.e. $H(dx) = 0$), F must assign no mass ($F(dx) = 0$). In the language of mathematical probability, F is said to be absolutely continuous with respect to H . To gain a better understanding of this condition, suppose that F and H both have densities that are positive everywhere. Then

$$k(x; s', e, s, e^*) = \frac{f(x; s', e, s, e^*)}{h(x; s', e, s, e^*)}$$

$$k(x, e) = f(x, e)/h(x, e).$$

However, F need not have a density in order that condition 2) hold. In particular, suppose that F assigns probability f_k to the point x_k for $k = 1, 2, \dots$. If H assigns mass h_k to the same sequence, then

$$k(x_k; s', e, s, e^*) = f_k/h_k.$$

Mixed probability distributions, that contain both point masses and densities, can be similarly handled.

We turn now to the application of importance sampling to the estimation problem (19). As argued previously, $f(X(t) : 0 \leq t < S_n) = \tilde{f}(A_0, \underline{C}_0, \dots, A_n, \underline{C}_n)$, so

$$\begin{aligned} \alpha &= E_F \tilde{f}(A_0, \underline{C}_0, \dots, A_n, \underline{C}_n) \\ &= \sum_{a_0, \dots, a_n} \int \cdots \int \tilde{f}(a_0, \underline{c}_0, \dots, a_n, \underline{c}_n) \mu(a_0) \prod_{e \in E(a_0)} F(d\underline{c}_0(e), e) \\ &\quad \cdot \prod_{j=0}^{n-1} p(a_{j+1}, ; a_j, e_j^*) \prod_{j=1}^n \prod_{e \in N(a_j, a_{j-1}, e_{j-1}^*)} F(d\underline{c}_j(e); a_j, e, a_{j-1}, e_{j-1}^*). \end{aligned}$$

By multiplying and dividing appropriately in the above expression (we need 1) and 2) above to avoid dividing by zero), we get

$$\begin{aligned} \alpha &= E_H \tilde{f}(A_0, \underline{C}_0, \dots, A_n, \underline{C}_n) L_n(F, H) \\ &= E_H f(X(t) : 0 \leq t < S_n) L_n(F, H) \end{aligned}$$

where

$$\begin{aligned} L_n(F, H) &= \frac{\mu(A_0)}{\gamma(A_0)} \prod_{e \in E(A_0)} k(C_0(e), e) \prod_{j=0}^{n-1} \frac{p(A_{j+1}; A_j, e_j^*)}{q(A_{j+1}; A_j, e_j^*)} \\ &\quad \cdot \prod_{j=1}^n \prod_{e \in N(A_j, A_{j-1}, e_{j-1}^*)} k(C_j(e); A_j, e, A_{j-1}, e_{j-1}^*) \end{aligned} \quad (20)$$

and $E_H(\cdot)$ is the expectation on the path space of X associated with the sampling distribution described above. The extension from deterministic n to stopping times T (i.e. $\{T = n\}$ is completely determined by $(A_0, \underline{C}_0, \dots, A_n, \underline{C}_n)$.) is straightforward, thereby yielding our next theorem.

Theorem 4. If $E_F |f(X(t) : 0 \leq t < S_T)| < \infty$ with T a (finite-valued) stopping time, then

$$E_F f(X(t) : 0 \leq t < S_T) = E_H f(X(t) : 0 \leq t < S_T) L_T(F, H)$$

where $L_T(F, H)$ is the r.v. obtained by substituting T into (20).

As in the continuous-time chain setting, the most important application of the stopping time version described above is to functions f of the form $f(X) = f(X(t) : 0 \leq t \leq w)$. Set $T(w) = \inf\{n \geq 0 : S_n > w\}$, and note that $T(w)$ is a stopping time. It follows that

$$E_F f(X(t) : 0 \leq t \leq w) = E_H f(X(t) : 0 \leq t \leq w) L_{T(w)}(F, H).$$

Thus, to apply importance sampling to computation of expectations of the form $\alpha = E_F f(X(t) : 0 \leq t \leq w)$, it suffices to simulate the GSMP under $P_H(\cdot)$ up to the $T(w)$ 'th transition. From the path so generated, one can easily recursively compute $L_{T(w)}(F, H)$, thereby yielding the observation $f(X(t) : 0 \leq t \leq w) L_{T(w)}(F, H)$. A sample mean consisting of independent replicates of the chain is then a consistent estimator of α .

We turn now to a discussion of the application of importance sampling to steady-state estimation of discrete-event simulations. To be precise, suppose that

$$f(X) = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t h(X(s)) ds,$$

where X is a GSMP. As discussed in Section 2 for discrete time chains, naive importance sampling breaks down, in the sense that if one sets $T = \infty$ in Theorem 4, the result becomes false. In fact, we have been able to obtain an importance sampling estimator for the steady-state mean of a GSMP only when the GSMP is regenerative.

To obtain regenerative structure, we assume that the GSMP hits a single state infinitely often, by which we mean that there exists $s_0, s_1 \in S$, $e_0 \in E(s_0)$ such that:

- a) $P_F\{(A_n, e_n^*, A_{n+1}) = (s_0, e_0, s_1) \text{ infinitely often}\} = 1$.
- b) $N(s_1, s_0, e_0) = E(s_1)$.

The term "single state" refers to the fact that if $E(s_0) = \{e_0\}$, then condition b) is automatically fulfilled. By b), all the clocks in $E(s_1)$ are reset, independently of the past, whenever s_1 is entered from s_0 by the triggering of event e_0 . If $\tau = \inf\{n \geq 1 : (A_{n-1}, e_{n-1}^*, A_n) = (s_0, e_0, s_1)\}$, it is then immediate that S_τ is a regeneration time for $X(\cdot)$. Note that the single-server queue is a GSMP with single set $(s_0 = 0, e_0 = \text{arrival event}, s_1 = 1)$, so that the regenerative structure of the $G/G/1$ queue follows as a special case of this discussion.

Letting $Y = \int_0^{S_\tau} h(X(s)) ds$, $\tilde{Y} = \int_0^{S_\tau} |h(X(s))| ds$, standard regenerative analysis shows that if $E_F \tilde{Y} < \infty$, then

$$E_F \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t h(X(s)) ds = \frac{E_H Y L_\tau(\underline{F}, \underline{H})}{E_H S_\tau L_\tau(\underline{F}, \underline{H})} \quad (22)$$

where $E_{\underline{H}}$ is the expectation on path space in which the initial distribution is given by $\nu(s_1) = 1$, $H(\cdot, e) = H(\cdot; s_1, e, s_0, e_0)$.

In fact, a second ratio formula for α involving the sampling distribution \underline{H} can be derived; this formula is the GSMP analogue of the discrete-time chain result obtained in Proposition 3. Observe that on the event $\{r > \ell\}$,

$$E_{\underline{H}} \left\{ \prod_{j=\ell}^{r-1} \frac{p(A_{j+1}; A_j, e_j^*)}{q(A_{j+1}; A_j, e_j^*)} \prod_{e \in N(A_{j+1}, A_j, e_j^*)} k(C_{j+1}(e); A_{j+1}, e, A_j, e_j^*) | A_0, C_0, \dots, A_\ell, C_\ell \right\} = 1$$

Since Y can be expressed as

$$\int_0^{S_r} h(X(s)) ds = \sum_{\ell=0}^{r-1} h(A_\ell) C_\ell(e_\ell^*) / r(A_\ell, e_\ell^*),$$

it follows, by conditioning the ℓ 'th summand of $Y L_r(\underline{F}, \underline{H})$ on $A_0, C_0, \dots, A_\ell, C_\ell$, that

$$E_{\underline{H}} \{Y L_r(\underline{F}, \underline{H})\} = E_{\underline{H}} \left\{ \sum_{\ell=0}^{r-1} h(A_\ell) C_\ell(e_\ell^*) L_\ell(\underline{F}, \underline{H}) / r(A_\ell, e_\ell^*) \right\}. \quad (23)$$

By setting $h = 1$ in (23), a similar formula for $E_{S_r} L_r(\underline{F}, \underline{H})$ is obtained. Thus, if $E_{\underline{F}}(\tilde{Y} + S_r) < \infty$, we find that

$$E_{\underline{F}} \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t h(X(s)) ds = \frac{E_{\underline{H}} \left\{ \sum_{\ell=0}^{r-1} h(A_\ell) C_\ell(e_\ell^*) L_\ell(\underline{F}, \underline{H}) / r(A_\ell, e_\ell^*) \right\}}{E_{\underline{H}} \left\{ \sum_{\ell=0}^{r-1} C_\ell(e_\ell^*) L_\ell(\underline{F}, \underline{H}) / r(A_\ell, e_\ell^*) \right\}} \quad (24)$$

As a consequence of the ratio formulas (22) and (24), it is clear that the steady-state mean α can be estimated by simulating replicates of either

$$(Y L_r(\underline{F}, \underline{H}), S_r L_r(\underline{F}, \underline{H}))$$

or

$$\left(\sum_{\ell=0}^{r-1} h(A_\ell) C_\ell(e_\ell^*) L_\ell(\underline{F}, \underline{H}) / r(A_\ell, e_\ell^*), \sum_{\ell=0}^{r-1} C_\ell(e_\ell^*) L_\ell(\underline{F}, \underline{H}) / r(A_\ell, e_\ell^*) \right)$$

under the sampling distribution $P_{\underline{H}}(\cdot)$. Let $\underline{r}_1(t)$, $\underline{r}_2(t)$ be the two corresponding ratio estimators that are available after t units of computational effort have been expended, and suppose $\underline{\beta}(1), \underline{\beta}(2)$ are r.v.'s denoting the computational effort required for each observational pair. By applying the methods of Section 5 of Glynn and Whitt (1986), the following central limit theorems follow.

Theorem 5. If $E_H(\beta(1) + \beta(2) + (Y^2 + S_r^2)L_r^2(\underline{F}, \underline{H})) < \infty$, then

$$t^{1/2}(\underline{x}_i(t) - \alpha) \Rightarrow \underline{\sigma}_i N(0, 1)$$

as $t \rightarrow \infty$, for $i = 1, 2$, where

$$\begin{aligned}\underline{\sigma}_1^2 &= E_H \beta(1) \cdot \frac{E_H\{(Y - \alpha S_r)^2 L_r^2(\underline{F}, \underline{H})\}}{(E_H S_r L_r(\underline{F}, \underline{H}))^2} \\ \underline{\sigma}_2^2 &= E_H \beta(2) \cdot E_H \left\{ \left(\sum_{\ell=0}^{r-1} (h(A_\ell) - \alpha) C_\ell(e_i^*) L_\ell(\underline{F}, \underline{H}) / r(A_\ell, e_i^*) \right)^2 \right\} / (E_H S_r L_r(\underline{F}, \underline{H}))^2.\end{aligned}$$

As in the discrete-time Markov chain case, the denominators of the above expressions for $\underline{\sigma}_1^2$ and $\underline{\sigma}_2^2$ are both equal to $(E_H S_r)^2$. Then, if $E_H \beta(1) \approx E_H \beta(2)$, $\underline{\sigma}_2^2$ is smaller than $\underline{\sigma}_1^2$ if the remaining expression in the numerator of $\underline{\sigma}_2^2$ is smaller than the corresponding expression in $\underline{\sigma}_1^2$. Again, as in the discrete-time chain case, we have no sufficient condition guaranteeing that $\underline{\sigma}_2^2 \leq \underline{\sigma}_1^2$.

We conclude this section by illustrating that the estimators of this section are generally easy to calculate. If $P_F(\cdot)$, $P_H(\cdot)$ correspond to $G/G/1$ queues (with identical initial distributions) having interarrival distributions F_A, H_A and service distributions F_S, H_S respectively, then

$$L_n(F, H) = \prod_{j \in \Lambda_n} k_A(C_j(1)) \prod_{j \in \Lambda_n^c} k_S(C_j(2))$$

where $\Lambda_n = \{j : j \leq n, S_j \text{ is an arrival epoch}\}$, $\Lambda_n^c = \{1, \dots, n\} - \Lambda_n$, and $k_A(\cdot)$, $k_S(\cdot)$ are the functions appearing in condition 2) of this section (k_A for the arrival event, k_S for the departure event). Hence,

$$L_n(F, H) = \begin{cases} L_{n-1}(F, H) k_A(C_n(1)), & n \in \Lambda_n \\ L_{n-1}(F, H) k_S(C_n(2)), & n \in \Lambda_n^c. \end{cases}$$

This simple recursive structure of $L_n(F, H)$ is typical of GSMP's.

5. IMPORTANCE SAMPLING: THE GENERAL THEORY

In the previous three sections, we have shown how to apply importance sampling to a broad class of stochastic simulations. The idea is to find a sampling distribution from which variates can be cheaply generated and for which the variance of the estimator is reduced. In this section, we abstract the ideas of the previous sections with the hope that the basic concepts underlying importance sampling will become more transparent.

Given a sample space Ω (with corresponding σ -field \mathcal{F}), suppose X is a real-valued r.v. defined on Ω for which we are required to calculate

$$\alpha = E_P X \quad (25)$$

where $E_P(\cdot)$ corresponds to the expectation on Ω associated with the probability distribution P . To define an appropriate sampling distribution Q for the estimation problem (25), we need an analogue to the conditions above given by (3) and (15) (as well as 1) and 2) of Section 4). It turns out that the abstract form of these conditions is:

$$\begin{aligned} P(A) > 0 \text{ implies } Q(A) > 0 \\ \text{for every } A \in \mathcal{F}. \end{aligned} \quad (26)$$

In the language of mathematical probability, (26) states that P is absolutely continuous with respect to Q . The Radon-Nikodym theorem then guarantees the existence of a "density" L_Q such that

$$P(A) = \int_A L_Q(\omega) Q(d\omega),$$

from which it follows that if $E_P|X| < \infty$,

$$E_P X = E_Q X L_Q. \quad (27)$$

L_Q is called the Radon-Nikodym derivative of P with respect to Q ; in the language of statistics, L is known as the likelihood ratio of P with respect to Q (hence, our use of the notation L). The derivative L_Q appears in all the main formulas of Sections 2 through 4 (e.g. (8), (16), (20)), in the form of the r.v.'s L_n, L_T , etc. appearing there.

Although condition (26) was the one used to construct sampling distributions in Sections 2 through 4, it turns out that importance sampling can be generalized somewhat beyond (26). Assume that Q satisfies:

$$\begin{aligned} E_P|X|I(A) > 0 \text{ implies } Q(A) > 0 \\ \text{for every } A \in \mathcal{F}. \end{aligned} \quad (28)$$

Since the set of A 's for which $E_P|X|I(A) > 0$ is contained in the set for which $P(A) > 0$, it follows that (28) is a weaker condition in Q than (26). By the Radon-Nikodym theorem, (28) suffices to guarantee existence of L_Q^* such that

$$E_P X I(A) = \int_A L_Q^*(\omega) Q(d\omega). \quad (29)$$

In this case, $\alpha = E_Q L_Q^*$. Here, L_Q^* has no interpretation as a likelihood ratio.

To compute the efficiency of an importance sampling estimator based on (29), we let β denote a r.v. representing the computation time required to calculate L_Q^* under sampling distribution Q . If $r(t)$ is the estimator available after t units of effort have been expended (see (5) and (6) for an explicit description), then $E_Q(\beta + L_Q^*) < \infty$ implies that

$$t^{1/2}(r(t) - \alpha) \Rightarrow \sigma(Q)N(0, 1)$$

where $\sigma^2(Q) = E_Q \beta \cdot \text{var}_Q L_Q^*$. Ideally, we would like to choose Q to minimize $\sigma^2(Q)$. This problem is difficult to solve mathematically since the quantity $E_Q \beta$ depends not only on the "complexity" of the sampling distribution Q but also on the efficiency of the programming implementation. As a result, most analyses of this problem disregard the $E_Q \beta$ term, and concentrate only on reducing the variance $\text{var}_Q L_Q^*$.

To minimize $\text{var}_Q L_Q^*$ subject to (29), observe that the Cauchy-Schwartz inequality yields

$$\begin{aligned} \text{var}_Q L_Q^* &= E_Q (L_Q^*)^2 - \alpha^2 \\ &\geq (E_Q |L_Q^*|)^2 - \alpha^2 \\ &= E_{\tilde{Q}} (L_Q)^2 - \alpha^2 = \text{var}_{\tilde{Q}} L_Q^* \end{aligned}$$

where $\tilde{Q}(A) = E_P |X| I(A) / E_P |X|$. Hence, \tilde{Q} is the sampling distribution which minimizes $\text{var}_Q L_Q^*$. This result is well known in the case where $\Omega = \mathcal{R}$ (see, for example, Rubinstein (1981), p. 122), where it is often pointed out that \tilde{Q} is not implementable practically since its construction requires knowledge of $E_P |X|$, which is generally unknown. The same proviso holds here, of course. Instead, one can exploit this result by choosing Q to minimize \tilde{Q} as closely as possible. Hence, the sampling distribution Q should inflate the probability mass assigned by P where $|X|$ is large, and deflate it where $|X|$ is small. Accomplishing this in practice is hard, however, and must generally be guided by sophisticated theory (see the "large deviations" applications of Section 7).

The form of the optimal \tilde{Q} does provide us with some limited information of general applicability. Specifically, $\tilde{Q}(A)$ is positive only when $E_P |X| I(A) > 0$. This implies that if $P(i, j) = 0$, then any Markov probability used for importance sampling should satisfy $K_\ell(i, j) = 0$.

Before proceeding, it is worth noting that (27) can be viewed in a somewhat different way. Given a sampling distribution Q , (27) shows that for every r.v. X defined on Ω and

every probability P which is absolutely continuous with respect to Q , (27) holds. Thus, by generating a sample point ω in Ω under Q , we simultaneously obtain an unbiased estimate of every expectation $\alpha = \alpha(X, P)$ of the form $\alpha(X, P) = E_P X$. For example, when one generates a path of a Markov chain having transition matrix P_1 , one is actually obtaining unbiased sampling information about every Markov chain having transition matrix P_2 for which $P_2(i, j) = 0$ whenever $P_1(i, j) = 0$.

We will now show that when Q is chosen to satisfy the stronger condition (26) (as occurs in Sections 2-4), importance sampling can be combined with the method of control variates to further reduce the variance. The main observation is that by setting $X = 1$ in (27), we find that $E_Q L = 1$. Hence, $Z(\lambda) = XL - \lambda(L - 1)$ has expectation α under Q . The basic theory of control variates then shows that if $E_Q(X^2 + 1)L^2 < \infty$, then $\text{var}_Q Z(\lambda)$ is minimized by choosing $\lambda = \lambda^*$, where

$$\lambda^* = \text{cov}_Q(XL, L) / \text{var}_Q L$$

(assuming $\text{var}_Q L > 0$), in which case $\text{var}_Q Z(\lambda^*) = \text{var}_Q XL - [\text{cov}_Q(XL, L)]^2 / \text{var}_Q L$. Of course, in a practical implementation, λ^* must generally be estimated from sample data.

One more point must be emphasized here. The use of importance sampling can frequently lead to estimators with infinite variance (see Proposition 6 for example); finite variance must not be taken for granted in this setting. This occurs despite the fact that importance sampling gives rise to estimators with finite first moment whenever the original estimation problem has finite expectation. Hence, if the sampling distribution Q is poorly chosen, importance sampling can lead to an estimation procedure with exceptionally poor convergence properties.

As discussed above, sampling under Q provides information about all probabilities P which satisfy (26). As we shall now see, Q can often also be used to generate sample outcomes in Ω having distribution P . Suppose the likelihood ratio L_Q is bounded by a (deterministic) constant M . If W is a random point in Ω generated under Q and if U is an independent uniform $(0, 1)$ r.v., then

$$\begin{aligned} Q(W \in d\omega | MU \leq L(W)) &= \frac{E_Q E_Q \{I(W \in d\omega, MU \leq L(W)) | W\}}{E_Q E_Q \{I(MU \leq L(W)) | W\}} \\ &= E_Q I(W \in d\omega) L(\omega) = P(d\omega). \end{aligned}$$

Thus, conditional on the event $\{MU \leq L(W)\}$, W has distribution P . This is the method of acceptance-rejection extended from $\Omega = \mathbb{R}$ to arbitrary Ω (see p. 50 of Rubinstein (1981)).

We can illustrate the technique with an application to the $M/M/1/\infty$ queue with time-varying arrival and service time rates. Specifically, we suppose that the arrival and service time rates at time t are given, respectively, by $\lambda(t)$ and $\mu(t)$ where $0 < \lambda(t), \mu(t) \leq K$. Our goal is to generate sample paths of an $M/M/1/\infty$ queue with these rates up to time S_n , where S_n is the entry instant of the $(n+1)$ 'st state visited.

If Q is the probability corresponding to a queue in which the rates are set equal to λ and μ for all $t \geq 0$, then Brémaud (1981), p. 168, shows that

$$L_Q = \prod_{j \in \Lambda_n} \left(\frac{\lambda(S_j)}{\lambda} \right) \prod_{j \in \Lambda_n^c} \left(\frac{\mu(S_j)}{\mu} \right) \exp \left[- \int_0^{S_n} (\lambda(s) - \lambda + (\mu(s) - \mu) I(Q(s) > 0)) ds \right]$$

where $Q(s)$ is the queue-length at time s , $\Lambda_n = \{j \leq n : S_j \text{ is an arrival epoch}\}$, and $\Lambda_n^c = \{1, \dots, n\} - \Lambda_n$. If λ, μ are lower bounds on the functions $\lambda(\cdot), \mu(\cdot)$ over $[0, t]$, it follows that $L_Q \leq K^{2n}/(\lambda\mu)^n$; the above acceptance-rejection method then applies to this problem. Hence, the time-varying queue previously described can be generated from trajectories of a "standard" $M/M/1/\infty$ queue. We should caution, however, that this would be an inefficient way of simulating a time-varying queue, since a (very) large proportion of the "standard" paths might be rejected. Competing methods for generation of sample paths for $M/M/1/\infty$ queues with time-varying arrival and service rates appear in Lewis and Shedler (1979 a, b).

6. IMPORTANCE SAMPLING: SECTIONS 2-4 REVISITED

We now wish to apply some of the ideas described in Section 5 to the processes discussed in Sections 2-4. We start by trying to determine the probability Q which minimizes $\text{var}_Q L_Q^*$ (see (29)), when P is a Markov probability associated with a discrete-time chain. Specifically, suppose our task is to estimate $\alpha = E_P f(X_0, \dots, X_n)$, where $E_P(\cdot)$ is the expectation on the path space of X associated with initial distribution μ and transition matrix P . It is easily verified that the probability Q is then given by

$$Q\{X_0 = i_0, \dots, X_n = i_n\} = \frac{|f(i_0, \dots, i_n)| \mu(i_0) \prod_{j=0}^{n-1} P(i_j, i_{j+1})}{E_P |f(X_0, \dots, X_n)|}. \quad (30)$$

Clearly, Q is not, in general, itself a Markov probability, in the sense that there will usually exist no initial distribution ν and transition matrices K_j for which $Q\{X_0 = i_0, \dots, X_n = i_n\} = \nu(i_0) \prod_{j=0}^{n-1} K_j(i_j, i_{j+1})$. Thus, the minimal variance probability Q has the additional disadvantage, beyond those discussed in Section 5, that Q will generally be non-Markov. The

non-Markov structure of \tilde{Q} suggests that generation of sample outcomes under \tilde{Q} will frequently be (very) computationally intensive. A similar conclusion holds for continuous-time chains and GSMP's.

There is one important setting, however, where the optimal measure \tilde{Q} is in fact Markov. Consider the problem of estimating, for a Markov chain X and subset A , the parameter

$$\alpha = P_P\{T(A) > m | X_0 = i\}$$

for $i \notin A$, where $T(A) = \inf\{n \geq 0 : X_n \in A\}$. Notice that $\alpha = E_P I(T(A) > m)$ where $I(T(A) > m)$ takes the multiplicative form

$$I(T(A) > m) = \prod_{k=0}^m I(X_k \notin A)$$

Then, (30) takes the form

$$\tilde{Q}\{X_0 = i_0, \dots, X_n = i_n\} = \delta_{i, i_0} \prod_{j=0}^{m-1} K_j(i_j, i_{j+1})$$

where

$$K_j(i, k) = P(i, k) \frac{P_P\{T(A) > m - j - 1 | X_0 = k\}}{P_P\{T(A) > m - j | X_0 = k\}}$$

for $i, k \notin A$.

For a birth-death process X in which $A = \{b+1, b+2, \dots\}$, it is evident that $P_P\{T(A) > m-1 | X_0 = i-1\} \geq P_P\{T(A) > m | X_0 = i\}$ so $K_j(i, i-1) \geq P(i, i-1)$ for all i and j . In other words the optimal measure \tilde{Q} increases the likelihood of drifting to the left. It is interesting to also note that $K_j(i, k) = P(i, k)$ if $|b-i| > m-j$.

Further information can be obtained by analyzing the asymptotic behavior of $P_P\{T(A) > m | X_0 = k\}$. Note that

$$P_P\{X_m = \ell, T(A) > m | X_0 = k\} = G_{k\ell}^m$$

where $G = (G_{k\ell} : k, \ell \in A^c)$ is the matrix in which $G_{k\ell} = P(k, \ell)$. Suppose X is irreducible, A^c has finite cardinality r , and G is diagonalizable. Then,

$$G = B^{-1} \text{diag}(\lambda_1, \dots, \lambda_r) B$$

is the spectral decomposition of G , where $|\lambda_i| < 1$ (this follows from the irreducibility). If $|\lambda_1| > |\lambda_2| > \dots > |\lambda_r|$, then

$$\lambda_1^{-m} G^m \rightarrow G^* = B^{-1} \text{diag}(1, 0, \dots, 0) B$$

(note that $\lambda_1 = |\lambda_1|$ by the Perron-Frobenius theorem for non-negative matrices). Thus, it follows that

$$G_{kl}^m \sim G_{kl}^* \cdot \lambda_1^m.$$

Under these conditions, one finds that $K_j(i, k) \rightarrow P(i, k) \cdot \lambda_1^{-1} \cdot \sum_{\ell \in A} G_{k\ell}^* / \sum_{\ell \in A} G_{i\ell}^*$. Hence, the optimal measure \tilde{Q} is asymptotically (for large m) time-homogeneous Markov, in the sense that $K_j(i, k) \rightarrow K(i, k)$ as $m \rightarrow \infty$ (for j fixed).

The above arguments suggest that it is frequently reasonable to reduce the variance by choosing a sampling distribution which minimizes $\text{var}_Q L_Q^*$ over the class of time-homogeneous Markov sampling distributions. Suppose that we wish to estimate

$$\alpha = E_P f(X_0, \dots, X_n)$$

where $X = \{X_n : n \geq 0\}$ is a discrete-time chain with transition matrix P and initial distribution μ . If Q is the Markov probability corresponding to transition matrix K and initial distribution ν , then

$$\begin{aligned} \text{var}_Q L_Q^* &= E_Q f^2(X_0, \dots, X_n) \frac{\mu(A_0)^2}{\nu(X_0)^2} \prod_{j=0}^{n-1} \frac{P(X_j, X_{j+1})^2}{K(X_j, X_{j+1})^2} - \alpha^2 \\ &= \sum_{i_0, i_1, \dots, i_n} h(i_0, \dots, i_n) \nu(i_0)^{-1} \cdot \prod_{j=0}^{n-1} K(i_j, i_{j+1})^{-1} - \alpha^2 \end{aligned}$$

where $h(i_0, \dots, i_n) = f^2(i_0, \dots, i_n) \mu^2(i_0) \prod_{j=0}^{n-1} P^2(i_j, i_{j+1})$. To find the minimizer K^*, ν^* for $\text{var}_Q L_Q^*$, we observe that the $n^2 + n$ variables $\{K(i, j), \nu(i) : 1 \leq i, j \leq n\}$ satisfy the constraints $K(i, j) \geq 0$, $\nu(i) \geq 0$, $\sum_{j=1}^n K(i, j) = 1$ ($1 \leq i \leq n$), $\sum_{i=1}^n \nu(i) = 1$. Applying the method of Lagrange multipliers, we find that the optimizer K^*, ν^* satisfies

$$\sum_{\substack{i_0, \dots, i_n \\ (i, j) \in \text{path}}} \frac{h(i_0, \dots, i_n)}{\prod_{(i_k, i_{k+1}) \neq (i, j)} K^*(i_k, i_{k+1})} \cdot \frac{\nu^*(i_0)^{-1}}{K^*(i, j)^{n(i_0, \dots, i_n)+1}} = \lambda_i, \quad 1 \leq i, j \leq n$$

$$\nu^*(i) = \tilde{\nu}(i) / \sum_{j=1}^n \tilde{\nu}(j), \quad 1 \leq i \leq n$$

where the λ_i 's are the Lagrange multipliers and

$$\begin{aligned} n(i_0, \dots, i_n) &= \text{total number of times } (i, j) \\ &\text{appears in path } (i_0, \dots, i_n) \end{aligned}$$

$$\tilde{\nu}(i) = \left\{ \sum_{i_1, \dots, i_n} \frac{h(i, i_1, \dots, i_n)}{K^*(i, i_1) \prod_{j=1}^{n-1} K^*(i_j, i_{j+1})} \right\}^{1/2}.$$

By multiplying through the (i, j) 'th equation by the corresponding denominator, we see that computing the $K^*(i, j)$'s involves solving a large number of coupled equations which are multinomial in the optimal probabilities $K^*(i, j)$. As a consequence, we conclude that it is, in general, unrealistic to expect that one can solve explicitly for the optimal Markovian sampling distribution Q^* .

A second important concept discussed in Section 5 is that the likelihood ratio L can be used as a control variate. In fact, for the processes of Section 2-4, control variates arise naturally when importance sampling is applied. To illustrate this idea, we focus on GSMP's. We first observe that

$$E_H \left\{ \frac{p(A_{n+1}; A_n, e_n^*)}{q(A_{n+1}; A_n, e_n^*)} \prod_{e \in N(A_{n+1}, A_n, e_n^*)} k(C_{n+1}(e); A_{n+1}, e, A_n, e_n^*) | A_0, C_0, \dots, A_n, C_n \right\} = 1;$$

this in turn implies that for $k \geq 0$,

$$E_H \{ L_{n+k}(F, H) | A_0, C_0, \dots, A_n, C_n \} = L_n(F, H). \quad (31)$$

(31) states that $\{L_n(F, H) : n \geq 0\}$ is a martingale sequence when generated under $P_H(\cdot)$. These simple observations guarantee that $E_H D_n(i) = 0$, where $D_n(i) = (L_{n+i}(F, H)/L_n(F, H)) - 1$. Hence, any of the r.v.'s $\{D_n(i) : n, i \geq 0\}$ may be used as controls.

Even more controls can be constructed when the estimation problem takes the form

$$\alpha = E_F \left\{ \int_0^{\tau} h(X(s)) ds \right\} \quad (32)$$

for some function h , with τ a stopping time. In this case, Section 4 (see (23), for example) shows that $E_H M = 0$, where

$$M = \sum_{\ell=0}^{\tau-1} (h(A_\ell) C_\ell(e_\ell^*) / r(A_\ell, e_\ell^*)) (L_\tau(F, H) - L_\ell(F, H)).$$

In fact, even more is true. Each of the summands appearing in M , namely

$$M_\ell = (h(A_\ell) C_\ell(e_\ell^*) / r(A_\ell, e_\ell^*)) (L_\tau(F, H) - L_\ell(F, H)) I(\tau > \ell)$$

have vanishing expectation with respect to E_H ; this follows immediately from the martingale property (31). Since the single state simulation strategy of Section 4 involves a ratio of expectations of the form (32) (for the denominator, set $h = 1$), we see that the controls M, M_0, M_1, \dots arise naturally when steady-state simulation for GSMP's is carried out.

It is to be expected that the control variates described here will yield at least a moderate increase in the efficiency of importance sampling methods for GSMP's.

7. IMPORTANCE SAMPLING: THE PARAMETRIC THEORY

With the exception of our brief discussion of the probability distribution Q in Section 5, we have concentrated primarily on describing the basic implementation of importance sampling in a stochastic simulation context, without saying much about how to actually choose the new sampling distribution Q . In this section, we briefly address this difficult problem. We shall assume that the goal is to estimate $\alpha = E_P X$, when the probability distribution P can be embedded in a parameterized family of distributions P_θ (i.e. $P = P_{\theta_0}$ for some θ_0 in the parameter set). Our goal is to select θ^* to maximize the efficiency of importance sampling. By the central limit theorem of Section 5, this is equivalent to requiring that we choose θ^* to minimize

$$\sigma^2(\theta) = E_\theta \beta \cdot \text{var}_\theta XL(\theta),$$

where $E_\theta(\cdot)$ is the expectation corresponding to P_θ and $L(\theta)$ is the likelihood ratio of P_{θ_0} with respect to P_θ .

To give some idea of the large variance reductions possible, consider the problem of estimating

$$\alpha = E_P f(X_0, \dots, X_{\tau(A)})$$

where $\tau(A) = \inf\{n \geq 0 : X_n \in A\}$ with $A = \{b+1, \dots\}$ and X is a Markov chain with transition matrix P . Suppose P takes the form

$$P(0, j) = \delta_{1j}$$

$$P(i, j) = p\delta_{i, i-1} + q\delta_{i, i+1} \text{ for } i \geq 1.$$

Such a P can be embedded in the parameterized family $P(\theta)$, where

$$P(\theta, 0, j) = \delta_{1j}$$

$$P(\theta, i, j) = \theta\delta_{i, i-1} + (1-\theta)\delta_{i, i+1} \text{ for } i \geq 1.$$

Note that $E_\theta \beta$ is independent of θ . Since the chain at time $\tau(A)$ has taken $b-i$ more steps to the right than steps to the left, it follows that

$$L(\theta) = \left(\frac{p}{\theta}\right)^{\frac{\tau(A)}{2} - (b - X_0)} \left(\frac{q}{1-\theta}\right)^{\frac{\tau(A)}{2}}.$$

Thus, if $\theta = q$, $L(q) = (q/p)^{b-X_0}$ and hence

$$\begin{aligned} & \text{var}_q \{f(X_0, \dots, X_{\tau(A)}) | X_0 = i\} \\ &= E_q \{f^2(X_0, \dots, X_{\tau(A)}) L^2(q) | X_0 = i\} - \alpha^2 \\ &= (q/p)^{b-i} E_P \{f^2(X_0, \dots, X_{\tau(A)}) | X_0 = i\} - \alpha^2. \end{aligned}$$

So, if $p \gg q$ or $b \gg i$, the potential computational savings obtained by using importance sampling with $\theta = q$ are enormous.

This example indicates that the additional effort required in finding an efficient θ at which to simulate can have a significant payback. One general approach for accomplishing this is to use stochastic approximation (e.g. the Kiefer-Wolfowitz method) to minimize $\sigma^2(\theta)$. Since $\sigma^2(\theta)$ may be infinite over a substantial subset of the parameter space (see Section 8), this can make a stochastic approximation algorithm highly unstable. A second general approach to selecting an efficient θ involves using the mathematical theory of probability to select a θ which will be close to θ^* ; the relevant mathematics needed is the theory of "large deviations". (See Billingsley (1979) for a brief description.) Consider, for example, the problem of estimating

$$\alpha = \alpha(w) \equiv P\{W > w\} \quad (33)$$

where W is the steady-state waiting time of a customer in a $GI/G/1/\infty$ queue having arrival and service time distributions F_A and F_S , respectively. Clearly, if w is large, this tail probability will be expensive to calculate to a reasonable degree of relative accuracy.

We now embed P in a parametric family. Assume that F_A, F_S have moment generating functions φ_A, φ_S that converge in a neighborhood of zero. Let

$$F_A(\theta, dx) = e^{\theta x} F_A(dx) / \varphi_A(\theta)$$

$$F_S(\theta, dx) = e^{\theta x} F_S(dx) / \varphi_S(\theta);$$

then $F_A(\theta), F_S(\theta)$ are the distribution functions of non-negative r.v.'s, and can be viewed as the inter-arrival and service time distribution for the P_θ system.

To select the best possible θ , let $\varphi(\theta) = \varphi_S(\theta)\varphi_A(-\theta)$. Note that φ is the moment generating function of the difference between a service time r.v. and an (independent) inter-arrival time r.v. Because we need the mean inter-arrival time to be greater than the mean service time for the queue to be stable and (33) to make sense, it follows that φ has a negative derivative at $\theta = 0$. Because of the convexity of φ , it is then evident that any root $\theta \neq 0$ of $\varphi(\theta) = 1$ must be positive and unique. Assume such a root θ_1 exists.

To apply these ideas to (33), we need to use the well-known fact that

$$\alpha = P \left\{ \max_{0 \leq k < \infty} T_k > w \right\} \quad (34)$$

where $T_0 = 0; T_n = \sum_{k=1}^n (S_k - A_k)$, and the A_k, S_k 's are independent copies of the arrival and service time r.v.'s. Note that (34) is equal to

$$\alpha = P\{\tau(w) < \infty\} \quad (35)$$

where $\tau(w) = \inf\{n \geq 0 : T_n > w\}$.

We now apply importance sampling to (35) by using inter-arrival and service time distributions $F_A(\theta_1), F_S(\theta_1)$ respectively, rather than F_A, F_S . Then

$$P\{\tau(w) < \infty\} = P_0\{\tau(w) < \infty\} = E_{\theta_1} I(\tau(w) < \infty) L_{\tau(w)}(\theta_1)$$

where $L_{\tau(w)}(\theta_1)$ is the likelihood ratio given by

$$L_{\tau(w)}(\theta_1) = \exp(-\theta_1 T_{\tau(w)}) \varphi(\theta_1)^{\tau(w)}.$$

It is easily verified that $E_{\theta_1}(S_1 - A_1) > 0$ and hence $\tau(w) < \infty$ occurs with probability one under P_{θ_1} . We conclude that

$$\alpha(w) = E_{\theta_1} L_{\tau(w)}(\theta_1). \quad (36)$$

Siegmund (1976) argues that importance sampling, based on P_{θ_1} as in (36), is asymptotically more efficient (as $w \rightarrow \infty$) than any other choice of θ when estimation is based on the representation (35). Asmussen (1985) shows that the choice θ_1 continues to be the optimal choice in heavy traffic conditions where $E_0 A \approx E_0 S$. The efficiency increases that are typical when this "large deviations" idea is applied can be enormous; Asmussen (1985) reports efficiency increases on the order of a factor of 3 to a factor of 400. Work on extension of these ideas to queueing networks is active; see, for example, Cottrell et al. (1983), Parekh and Walrand (1986), Weiss (1986), and Anantharun (1987). This area of work holds great promise for development of importance sampling methods permitting substantial efficiency increases.

8. RESPONSE SURFACE ESTIMATION USING IMPORTANCE SAMPLING

In many applications, a parametric embedding of the type described in Section 7 arises naturally. For example, in the study of queueing systems, one is frequently interested in understanding the behavior of the process as a function of the probability law P_θ imposed on the queue, where P_θ corresponds to the queue dynamics under service rate θ . Since the family $\{P_\theta\}$ so obtained often satisfies (26) (i.e. $P_{\theta_0}(A) > 0$ implies $P_{\theta_1}(A) > 0$ for $\theta_0 \neq \theta_1$), importance sampling seems natural in this context. In particular, one can then estimate the expectation of X under P_θ by simulating X under P_{θ_0} for some fixed θ_0 . This suggests that one can use importance sampling, under distribution P_{θ_0} , to simultaneously estimate the entire response function

$$\alpha(\theta) = E_{P_\theta} X.$$

In a queueing context, this idea raises the possibility that one can estimate queue performance over an entire interval of service time parameters θ by simulating the queue at a particular θ_0 and applying importance sampling appropriately.

To describe the idea precisely, let X be a real-valued r.v. defined on a sample space Ω (with σ -field \mathcal{F}). Suppose $\{P_\theta : \theta \in [a, b]\}$ is a family of probabilities on (Ω, \mathcal{F}) which satisfy:

$$\begin{aligned} P_\theta(A) > 0 \text{ implies } Q(A) > 0 \\ \text{for every } A \in \mathcal{F} \text{ and } \theta \in [a, b]. \end{aligned} \quad (37)$$

In many applications, (37) is satisfied with $Q = P_{\theta_0}$, for some $\theta_0 \in [a, b]$. Then, it follows from Section 5 that for every $\theta \in [a, b]$, there exists $L(\theta)$ such that

$$E_{P_\theta} X = E_Q X L(\theta).$$

Hence, if $E_{P_\theta} |X| < \infty$ for all $\theta \in [a, b]$, it follows that

$$\alpha(\theta) = E_Q X L(\theta). \quad (38)$$

To estimate the "response function" $\alpha(\cdot)$, (38) suggests that one generate independent replicates $\{X_i L_i(\theta) : \theta \in [a, b]\}$ of the random process $\{X L(\theta) : \theta \in [a, b]\}$. The function $\alpha(\cdot)$ can then be estimated via

$$\bar{\alpha}_n(\cdot) = \frac{1}{n} \sum_{k=1}^n X_k L_k(\cdot). \quad (39)$$

To illustrate this idea, we now specialize to the discrete-time Markov chain setting. For each $\theta \in [a, b]$, let $P(\theta)$, $\mu(\theta)$ be an associated transition matrix and initial distribution, respectively; the parameter θ might correspond, for example, to a branching probability, in which case $P(\theta)$ is the transition matrix associated with branching probability θ . Suppose that the goal is to estimate the expectation of $f(X_0, X_1, \dots, X_n)$ under transition matrix $P(\theta)$ and initial distribution $\mu(\theta)$, as a function of θ . If $\Lambda(\theta) = \{(i, j, k) : P(\theta, i, j) > 0, \mu(\theta, k) > 0\}$ is independent of θ , then we can estimate $\alpha(\cdot)$ via

$$\bar{\alpha}_n(\cdot) = \frac{1}{n} \sum_{k=1}^n f(X_{k0}, X_{k1}, \dots, X_{kn}) L_{nk}(P(\cdot), P(\theta_0)) \quad (40)$$

where

$$L_{nk}(P(\theta), P(\theta_0)) = \frac{\mu(\theta, X_{k0})}{\mu(\theta_0, X_{k0})} \prod_{j=0}^{n-1} \frac{P(\theta, X_{kj}, X_{k,j+1})}{P(\theta_0, X_{kj}, X_{k,j+1})}$$

and $\{(X_{k0}, X_{k1}, \dots, X_{kn}) : k \geq 0\}$ is a sequence of i.i.d. copies of (X_0, X_1, \dots, X_n) generated under P_{θ_0} . The attractive feature of this estimator is that the global behavior of $\alpha(\cdot)$

may be estimated by simulating the Markov chain locally (at θ_0), so that one avoids the enormous computation involved in simulating the system at a grid of points $\theta_1, \theta_2, \dots, \theta_m$.

While it is true that $\bar{\alpha}_n(\theta) \rightarrow \alpha(\theta)$ a.s. for each $\theta \in [a, b]$ if $E_{P_\theta}|X| < \infty$, one is generally more interested, in the current function estimation context, in studying

$$\epsilon_n(a, b) \equiv \sup_{\theta \in [a, b]} |\bar{\alpha}_n(\theta) - \alpha(\theta)|.$$

The quantity $\epsilon_n(a, b)$ describes the global behavior of $\bar{\alpha}_n(\cdot)$ as an estimator of the function $\alpha(\cdot)$ over the interval (a, b) .

Theorem 6. Suppose that $L(\theta)$ is continuous in θ a.s. on $[a, b]$, so that $M = M(a, b) \equiv \max\{L(\theta) : \theta \in [a, b]\} < \infty$. If $E|X| \cdot M < \infty$, then $\epsilon_n(a, b) \rightarrow 0$ a.s. as $n \rightarrow \infty$.

This result, which states that $\bar{\alpha}_n(\cdot)$ approximates $\alpha(\cdot)$ uniformly well over the interval $[a, b]$, is an immediate consequence of Theorem 8.1, p. 256, of Parthasarathy (1967).

Returning to the Markov chain setting described above, assume that $P(\theta), \mu(\theta)$ are continuous in θ . Since f is automatically bounded (S is finite), it follows that $E|f(X_0, X_1, \dots, X_n)| \leq M < \infty$ for every $n \geq 0$ and all intervals $[a, b]$. Hence, $\bar{\alpha}_n(\cdot) \rightarrow \alpha(\cdot)$ uniformly on bounded intervals a.s.

The situation is more complicated, however, when steady-state estimation is involved. As discussed earlier, we suggest using regenerative simulation in this case (see Proposition 2).

As discussed in Proposition 2, it is evident that

$$\alpha(\cdot) = \frac{E_{P(\theta_0)}\{Y L_r(P(\cdot), P(\theta_0)) | X_0 = i\}}{E_{P(\theta_0)}\{\tau L_r(P(\cdot), P(\theta_0)) | X_0 = i\}} \equiv \frac{u(\cdot)}{\ell(\cdot)}$$

To estimate $\alpha(\cdot)$ uniformly well over an interval $[a, b]$, it suffices to uniformly estimate $u(\cdot), \ell(\cdot)$ over the interval. To apply Theorem 6 to $u(\cdot)$, observe that

$$|Y| \cdot \max\{L_r(P(\theta), P(\theta_0)) : \theta \in [a, b]\} \leq \|h\| \cdot \tau K(a, b)^\tau \quad (41)$$

where $\|h\| = \max\{|h(i)| : i \in S\}$ (h is the steady-state functional under consideration), $K(a, b) = \max\{P(\theta, i, j)/P(\theta_0, i, j) : P(\theta_0, i, j) > 0, \theta \in [a, b]\}$. Note that the expectation of the right-hand side of (41) is just $K_i(K(a, b))\|h\|$, where $K_j(z) = E_{P(\theta_0)}\{\tau z^\tau | X_0 = j\}$.

Proposition 6. If $0 < \sum_{k \neq i} P(\theta_0, j, k)$ for all $j \in S$ (S finite), then there exists z_0 such that $k_i(z) = \infty$ for $z > z_0$.

Thus, if $K(a, b) > z_0$, the bound (41) becomes infinite. This suggests that if $K(a, b)$ is large, $\hat{\alpha}_m(\cdot)$ may not approximate $\alpha(\cdot)$ uniformly well over $[a, b]$. Of course, $K(a, b)$ will tend to be large, for most parameterizations $P(\theta)$, when $b - a$ is big. Hence, steady-state response function estimation appears to behave poorly when the interval $[a, b]$ under consideration is too large.

This conclusion is also suggested by the following result. For the proof, see the next section.

Proposition 7. If $\inf\{\sum_{k \neq i} \frac{P^2(\theta, j, k)}{P(\theta_0, j, k)} : j \in S\} > 1$, and $\frac{P^2(\theta, j, 0)}{P(\theta_0, j, 0)} > 0$ for all j , then $\text{var}_{P(\theta_0)} L_r(P(\theta), P(\theta_0)) = \infty$.

9. COMPUTATIONAL AND EMPIRICAL RESULTS

Whenever new simulation methodology is proposed, it is always helpful to be able to find the theoretical values for quantities being estimated. In this section we shall discuss the relevant theoretical values in the context of the steady-state estimation problem for discrete-time Markov chains.

Suppose that $X = \{X_n : n \geq 0\}$ is an irreducible finite-state Markov chain with transition probability matrix $P = \{P(i, j)\}$ and state space $E = \{0, 1, \dots, N\}$. Assume we wish to estimate

$$\alpha = E_P \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} h(X_k)$$

for some given $h : E \rightarrow \mathbb{R}$. As indicated in Section 2, this problem can be attacked using the regenerative method. We shall use the state 0 as a "return" state in the regenerative method and set

$$\tau = \inf\{n \geq 1 : X_n = 0\} \quad \text{and} \\ Y = \sum_{k=0}^{\tau-1} h(X_k).$$

Standard regenerative theory yields $\alpha = E_P\{Y|X_0 = 0\}/E_P\{\tau|X_0 = 0\}$. If we now switch to a new transition matrix K (satisfying the conditions below equation (7)), we see from Proposition 2, that we also have

$$\alpha = \frac{E_K\{Y L_r | X_0 = 0\}}{E_K\{\tau L_r | X_0 = 0\}},$$

where $L_r = L_r(P, K)$ as in Proposition 1. From Proposition 1 we know that

$$E_P\{Y|X_0 = 0\} = E_K\{Y L_r | X_0 = 0\} \quad \text{and} \\ E_P\{\tau|X_0 = 0\} = E_K\{\tau L_r | X_0 = 0\}.$$

Let P denote the matrix obtained from P by replacing $P(0, j)$ by 0 for all $j \in E$. Then from standard theory (Hordijk, Iglehart, Schassberger (1976)), we know that

$$E_P\{Y|X_0 = 0\} = [(I - {}_0P)^{-1}h]_0 \quad \text{and}$$

$$E_P\{\tau|X_0 = 0\} = [(I - {}_0P)^{-1}h]_0,$$

where ϵ is the vector of all ones, h is a vector with components $h(i)$ ($i \in E$), and $[\]_0$ denotes the 0 component of the vector.

To compare the efficiency of estimating α under K versus P , we use Theorem 2. Assume that the computational effort to generate a cycle is proportional to the length of the cycle, so $\beta(i) = \tau$. The variance of the limiting normal when α is estimated under P is given by

$$\frac{E_P\{(Y - \alpha\tau)^2|X_0 = 0\}}{E_P\{\tau|X_0 = 0\}} \quad (42)$$

whereas if α is estimated under K , using the ratio representation given by Proposition 2, then the appropriate variance is

$$\frac{E_K\{(Y - \alpha\tau)^2 L_r^2(P, K)|X_0 = 0\} \cdot E_K\{\tau|X_0 = 0\}}{(E_K\{\tau L_r(P, K)|X_0 = 0\})^2} \quad (43)$$

To calculate the four terms appearing in (42) and (43), only the term $E_K\{(Y - \alpha\tau)^2 \cdot L_r^2(P, K)|X_0 = 0\}$ presents something new. (Recall that $E_K\{\tau L_r(P, K)|X_0 = 0\} = E_P\{\tau|X_0 = 0\}$.) The other three terms can be handled using the methods in Hordijk, Iglehart, Schassberger (1976). Calculating this new term is a formidable job and in fact it may be equal to $+\infty$.

We begin by computing the quantity

$$a_j = E_K\{L_r^2(P, K)|X_0 = j\}.$$

Note that

$$\begin{aligned} a_j &= \sum_{\ell=1}^{\infty} E_K\{L_r^2(P, K)I(\tau = \ell)|X_0 = j\} \\ &= \sum_{\ell=1}^{\infty} \sum_{\substack{i_0, \dots, i_{\ell} \\ i_0=j, i_{\ell}=0 \\ i_j \neq 0, 1 \leq j < \ell}} \frac{P^2(i_0, i_1)}{K(i_0, i_1)} \dots \frac{P^2(i_{\ell-1}, i_{\ell})}{K(i_{\ell-1}, i_{\ell})} \\ &= \sum_{\ell=1}^{\infty} (C^{\ell-1}b)_j \end{aligned}$$

where $C = (C_{k\ell})$, $b = (b_{\ell})$ and $C_{k\ell} = P^2(k, \ell)/K(k, \ell)$ for $\ell \neq 0$, $C_{k0} = 0$, $b_{\ell} = P^2(\ell, 0)/K(\ell, 0)$. Hence, we can compute the a_j 's by solving for

$$\alpha = \sum_{r=0}^{\infty} C^r b.$$

Note that a satisfies $a = Ca + b$. If a is finite, then $(I - C)a = b$, so that if $(I - C)^{-1}$ exists, a must equal $(I - C)^{-1}b$. Thus, even when $(I - C)^{-1}$ exists, we are still faced with the problem of distinguishing when a is finite, so that $a = (I - C)^{-1}b$.

Proposition 8. Assume $C \geq 0$, $a = \sum_{r=0}^{\infty} C^r b$, and $(I - C)^{-1}$ exists. If $(I - C)^{-1} \geq 0$, then a is finite and equals $(I - C)^{-1}b$. Furthermore, $(I - C)^{-1} \geq 0$ if and only if the spectral radius of C is less than one.

A sufficient condition for the spectral radius of C to be less than one is that

$$\max_{j \in E} \sum_{k=0}^N C_{jk} < 1$$

or, in our setting

$$\max_{j \in E} \sum_{k \neq 0} \frac{P^2(j, k)}{K(j, k)} < 1.$$

Similarly, it is easily verified that a sufficient condition for $\sum_{r=0}^{\infty} C^r = \infty$ is that

$$\min_{j \in E} \sum_{k=0}^N C_{jk} > 1$$

or

$$\min_{j \in E} \sum_{k \neq 0} \frac{P^2(j, k)}{K(j, k)} > 1.$$

This observation basically proves Proposition 7.

To compute the variance term (43), we need to compute the term

$$d_j = E_K \left\{ \left(\sum_{k=0}^{r-1} g(X_k) \right)^2 L_r^2(P, K) | X_0 = j \right\} \quad (44)$$

where $g(\cdot) = h(\cdot) - \alpha$. Expanding the square in (44) and conditioning on the value of X_1 leads to

$$\begin{aligned} d_j &= g^2(j) E_K \{ L_r^2(P, K) | X_0 = j \} \\ &+ 2g(j) \sum_{k \neq 0} \frac{P^2(j, k)}{K(j, k)} E_K \left\{ \sum_{r=0}^{r-1} g(X_r) L_r^2(P, K) | X_0 = k \right\} \\ &+ \sum_{k \neq 0} \frac{P^2(j, k)}{K(j, k)} d_k. \end{aligned}$$

So, d satisfies $d = b' + Cd$, where $C_{k\ell} = P^2(k, \ell)/K(k, \ell)$ for $\ell \neq 0$, $C_{K0} = 0$, and

$$\begin{aligned} b'_j &= g^2(j) E_K \{ L_r^2(P, K) | X_0 = j \} \\ &+ 2g(j) \sum_{k \neq 0} \frac{P^2(j, k)}{K(j, k)} E_K \left\{ \sum_{r=0}^{r-1} g(X_r) L_r^2(P, K) | X_0 = k \right\}. \end{aligned}$$

This system of linear equations has the same structure as that derived for the $E_K\{L_r^2(P, K)|X_0 = j\}$'s discussed above. Finally, an analogous system of equations can be obtained for the unknown term

$$E_K \left\{ \sum_{r=0}^{r-1} g(X_r) L_r^2(P, K) | X_0 = k \right\}$$

appearing in b_j' . The equations for the d_j 's can then be solved to calculate the variance term (43).

At this point we have only developed one numerical illustration of importance sampling for discrete-time Markov chains.¹ The original P matrix (corresponding to an (s, S) inventory model) is given by

$$P = \begin{bmatrix} 1/3 & 0 & 0 & 2/3 \\ 1/3 & 1/3 & 0 & 1/3 \\ 1/3 & 1/3 & 1/3 & 0 \\ 0 & 1/3 & 1/3 & 1/3 \end{bmatrix}$$

with state-space $E = \{1, 2, 3, 4\}$. We were interested in simulating to estimate $\Pi_4 = 0.347838$. The K matrix used in the importance sampling is given by

$$K = \begin{bmatrix} 0.476 & 0 & 0 & 0.524 \\ 0.311 & 0.476 & 0 & 0.213 \\ 0.206 & 0.319 & 0.475 & 0 \\ 0 & 0.233 & 0.233 & 0.534 \end{bmatrix}$$

Based on a simulation of 3000 4-cycles, we found an estimate of 0.319 (0.540) for the term $E_P\{(Y - \alpha\tau)^2 | X_0 = 4\}$ ($E_K\{(Y L_r - \alpha\tau L_r)^2 | X_0 = 4\}$). We can compute $E_K\{\tau | X_0 = 4\}$ ($E_P\{\tau | X_0 = 4\}$) to be 2.8749 (3.0625). This leads to a value for the ratio, R^2 , of (43) to (42) of 0.745. Using the alternate method proposed in Proposition 3 leads to a value for the corresponding R^2 ratio of 1.278. As indicated in Section 2, the asymptotic efficiency of importance sampling does depend on the return state used. For this example, the R^2 values for the first method using states 1, 2, and 3 respectively were 1.615, 1.568, and 2.153. For the second method the corresponding values were 1.822, 1.729, and 2.069. The fact that these values were larger than 1 should not be surprising, since the K matrix was specifically chosen to do well for the return state 4.

In conclusion, we need to obtain much more computational experience with importance sampling. In particular better methods for selecting the Q matrix are needed.

¹ We are grateful to Scott Schulz for providing this numerical example.

APPENDIX

Proof of Proposition 1: Note that

$$E_P f(X_0, \dots, X_T) = \sum_{k=0}^{\infty} E_P f(X_0, \dots, X_k) I(T = k) \quad (A1)$$

Since T is a stopping time, $f(X_0, \dots, X_k) I(T = k) = \tilde{f}_k(X_0, \dots, X_k)$ for some function \tilde{f}_k . By (8),

$$E_P \tilde{f}_k(X_0, \dots, X_k) = E_k \tilde{f}_k(X_0, \dots, X_k) L_k(P, K) \quad (A2)$$

Combining (A1) and (A2), we obtain

$$\begin{aligned} E_P f(X_0, \dots, X_T) &= \sum_{n=0}^{\infty} E_K f(X_0, \dots, X_n) I(T = n) L_n(P, K) \\ &= E_K f(X_0, \dots, X_T) L_T(P, K). \end{aligned}$$

Proof of Proposition 6. Let $\varphi_j(z) = E_{P(\theta_0)}\{z^j | X_0 = j\}$. By conditioning on X_1 , we find that

$$\varphi_j(z) = z P(\theta_0, j, i) + z \sum_{k \neq i} P(\theta_0, j, k) \varphi_k(z)$$

i.e. $\varphi(z) = a(z) + C(z)\varphi(z)$. Since $C(z), \varphi(z), a(z) \geq 0$ for $z > 0$, it follows that $\varphi(z) \geq \sum_{r=0}^{\infty} C^r(z)a(z)$. Also, for $z > z_0 \equiv 1/\min_j \sum_{k \neq i} P(\theta_0, j, k)$,

$$\min_j \sum_k C_{jk}(z) > 1$$

which implies that $C(z)^k \rightarrow \infty$ as $k \rightarrow \infty$. Since $a(z) > 0$, it follows that $\varphi(z) = \infty$ for $z > z_0$.

Proof of Proposition 8. Since $C \geq 0$, it is evident that $\sum_{j=0}^n C^j \cdot (I - C) = I - C^{n+1} \leq I$. Since $(I - C)^{-1} \geq 0$, $\sum_{j=0}^n C^j \leq (I - C)^{-1}$. By letting $n \rightarrow \infty$, it follows that $\sum_{j=0}^{\infty} C^j \leq (I - C)^{-1}$. Hence, $a < \infty$ and (as argued in Section 9) $a = (I - C)^{-1}b$.

This argument also proves that if $(I - C)^{-1} \geq 0$, then $C^j \rightarrow 0$ and $j \rightarrow \infty$; the fact that $C^j \rightarrow 0$ is equivalent to the spectral radius of C being less than one. For the converse, note that if $C^j \rightarrow 0$, it must necessarily tend to zero exponentially fast, so that $0 \leq \sum_{j=0}^{\infty} C^j < \infty$. But it is easily verified that $\sum_{j=0}^{\infty} C^j (I - C) = I$ so that $(I - C)^{-1} = \sum_{j=0}^{\infty} C^j < \infty$.

REFERENCES

- Anantharum, V. "A Transient Performance Criterion for Networks and its Asymptotically Optimal Fast Simulation for Tandems", forthcoming technical report, School of Electrical Engineering, Cornell University, Ithaca, NY, 1987.
- Asmussen, S. "Conjugate Processes and the Simulation of Ruin Problems", *Stoch. Proc. Appl.* **20**, 213-229, 1985.
- Billingsley, P. *Probability and Measure*, John Wiley, New York, 1979.
- Br  nard, P. *Point Processes and Queues: Martingales Dynamics*, Springer-Verlag, New York, 1981.
-   nlar, E. *Introduction to Stochastic Processes*, Prentice-Hall, Englewood Cliffs, NJ, 1975.
- Cottrell, M., Fort, J., and Malgouyres, G. "Large Deviations and Rare Events in the Study of Stochastic Algorithms", *IEEE Trans. on Aut. Cont.* **28**, 907-920, 1983.
- Crane, M. and Iglehart, D. L. "Simulating Stable Stochastic Systems, III: Regenerative Processes and Discrete-Event Simulations", *Oper. Res.*, **23**, 33-45, 1975.
- Fox, B. L. and Glynn, P. W. "Discrete-Time Conversion for Simulating Semi-Markov Processes", *Oper. Res. Letters*, **5**, 191-196 (1986).
- Glynn, P. W. "On the Role of Generalized Semi-Markov Processes in Simulation Output Analysis", *Proc. of the 1983 Winter Simulation Conference*, 38-42, 1983.
- Glynn, P. W. and Whitt, W. "The Efficiency of Simulation Estimators", submitted for publication, 1986.
- Hordijk, A., Iglehart, D. L., and Schassberger, R. "Discrete Time Methods for Simulating Continuous Time Markov Chains", *Adv. Appl. Prob.* **8**, 772-788, 1976.
- Lewis, P. A. W. and Shedler, G. A. "Simulation of Nonhomogeneous Poisson Processes by Thinning", *Nav. Res. Logistics Quart.* **26**, 403-414, 1979.
- Lewis, P. A. W. and Shedler, G. S. "Simulation of Nonhomogeneous Poisson Processes with Degree-two Exponential Polynomial Rate Function", *Oper. Res.* **27**, 1026-1040, 1979.
- Parekh, S. and Walrand, J. "Quick Simulation Method for Excessive Backlogs in Networks of Queues", Technical Report, Electronics Research Laboratory, University of California, Berkeley, CA, 1986.
- Parthasarathy, K. R. *Probability Measures on Metric Spaces*, Academic Press, New York, 1967.
- Rubinstein, R. Y. *Simulation and the Monte Carlo Method*, John Wiley, New York, 1981.

- Siegmund, D. "Importance Sampling in the Monte Carlo Study of Sequential Tests", *Ann. Statist.* **4**, 673-684, 1976.
- Weiss, A. "A New Technique for Analyzing Large Traffic Systems", *Adv. Appl. Prob.* **18**, 506-532, 1986.

A193 585

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER Technical Report No. 24	2. GOVT ACCESSION NO. N/A	3. RECIPIENT'S CATALOG NUMBER N/A
4. TITLE (and Subtitle) IMPORTANCE SAMPLING FOR STOCHASTIC SIMULATIONS		5. TYPE OF REPORT & PERIOD COVERED Technical Report
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Peter W. Glynn and Donald L. Iglehart		8. CONTRACT OR GRANT NUMBER(s) DAAG29-84-K-0030
9. PERFORMING ORGANIZATION NAME AND ADDRESS Department of Operations Research Stanford University Stanford, CA 94305-4022		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS N/A
11. CONTROLLING OFFICE NAME AND ADDRESS U.S. Army Research Office Post Office Box 12211 Research Triangle Park, NC 27709		12. REPORT DATE August 1987
		13. NUMBER OF PAGES 36 pp.
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report) NA		
18. SUPPLEMENTARY NOTES The view, opinions, and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Importance sampling, variance reduction, simulation, Monte Carlo method, Markov chains, generalized semi-Markov processes.		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) PLEASE SEE OTHER SIDE		

20. (Continued.)

Importance sampling is one of the classical variance reduction techniques for increasing the efficiency of Monte Carlo algorithms for estimating integrals. The basic idea is to replace the original random mechanism in the simulation by a new one and at the same time modify the function being integrated. In this paper the idea is extended to problems arising in the simulation of stochastic systems. Discrete-time Markov chains, continuous-time Markov chains, and generalized semi-Markov processes are covered. Applications are given to a $GI/G/1$ queueing problem and response surface estimation. Computation of the theoretical moments arising in importance sampling is discussed and some numerical examples given.

UNCLASSIFIED

FILMED
78